



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : H04L 12/46, H04Q 11/04	A1	(11) International Publication Number: WO 95/20282 (43) International Publication Date: 27 July 1995 (27.07.95)
(21) International Application Number: PCT/CA95/00029 (22) International Filing Date: 20 January 1995 (20.01.95) (30) Priority Data: 9401092.3 21 January 1994 (21.01.94) GB (71) Applicant (for all designated States except US): NEWBRIDGE NETWORKS CORPORATION [CA/CA]; 600 March Road, P.O. Box 13600, Kanata, Ontario K2K 2E6 (CA). (72) Inventors; and (75) Inventors/Applicants (for US only): BURWELL, Wayne [CA/CA]; 505 Lisgar Street, Ottawa, Ontario K1R 5H2 (CA). COOMBER, Dave [CA/CA]; 35 Horner Drive, Nepean, Ontario K2H 5E8 (CA). DUXBURY, Tom [CA/CA]; 530 Berry Side Road, Dunrobin, Ontario K0A 1T0 (CA). BEHKI, Nutan [CA/CA]; 40 Moorcroft Road, Nepean, Ontario K2G 0M7 (CA). GALWAY, Keith [CA/CA]; 1092 Bravar Drive, Manotick, Ontario K4M 1G3 (CA). WATT, James [CA/CA]; 1003-960 Teron Road, Kanata, Ontario K2K 2B6 (CA). GLENDINNING, Duncan [CA/CA]; 21 Tamblyn Crescent, Kanata, Ontario K2L 3A5 (CA). ZYWICKI, Eugene [CA/CA]; 9 Ardell Grove, Nepean, Ontario K2G 4G5 (CA).		(74) Agent: MITCHELL, Richard, J.; Marks & Clerk, P.O. Box 957, Station B, Ottawa, Ontario K1P 5S7 (CA). (81) Designated States: AM, AT, AU, BB, BG, BR, BY, CA, CH, CN, CZ, DE, DK, EE, ES, FI, GB, GE, HU, JP, KE, KG, KP, KR, KZ, LK, LR, LT, LU, LV, MD, MG, MN, MW, MX, NL, NO, NZ, PL, PT, RO, RU, SD, SE, SI, SK, TJ, TT, UA, US, UZ, VN, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG), ARIPO patent (KE, MW, SD, SZ). Published <i>With international search report.</i>
(54) Title: TRANSPARENT INTERCONNECTOR OF LANs BY AN ATM NETWORK (57) Abstract <p>A computer communications network comprises a plurality of interconnected ATM switches (1a, 1b, 1c) forming a WAN or LAN area network over which ATM cells are transmitted, and a plurality of devices (8, 9, 10, 11, 14) including LAN interface adapters for connection to one or more local area networks (LANs). Interface devices (6a, 6b, 6c) connect at least some of the respective ATM switches to the LAN interface adapters. The interface adapts the LAN frames for transport over the ATM network. The user devices can thus communicate through the LAN interface adapters transparently with the wide area network.</p>		

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	GB	United Kingdom	MR	Mauritania
AU	Australia	GE	Georgia	MW	Malawi
BB	Barbados	GN	Guinea	NE	Niger
BE	Belgium	GR	Greece	NL	Netherlands
BF	Burkina Faso	HU	Hungary	NO	Norway
BG	Bulgaria	IE	Ireland	NZ	New Zealand
BJ	Benin	IT	Italy	PL	Poland
BR	Brazil	JP	Japan	PT	Portugal
BY	Belarus	KE	Kenya	RO	Romania
CA	Canada	KG	Kyrgyzstan	RU	Russian Federation
CF	Central African Republic	KP	Democratic People's Republic of Korea	SD	Sudan
CG	Congo	KR	Republic of Korea	SE	Sweden
CH	Switzerland	KZ	Kazakhstan	SI	Slovenia
CI	Côte d'Ivoire	LI	Liechtenstein	SK	Slovakia
CM	Cameroon	LK	Sri Lanka	SN	Senegal
CN	China	LU	Luxembourg	TD	Chad
CS	Czechoslovakia	LV	Latvia	TG	Togo
CZ	Czech Republic	MC	Monaco	TJ	Tajikistan
DE	Germany	MD	Republic of Moldova	TT	Trinidad and Tobago
DK	Denmark	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	US	United States of America
FI	Finland	MN	Mongolia	UZ	Uzbekistan
FR	France			VN	Viet Nam
GA	Gabon				

TRANSPARENT INTERCONNECTOR OF LANs BY AN ATM NETWORK

This invention relates to a computer communications network.

There are various types of computer communications
5 networks for establishing communications between digital
user devices, such as computers and workstations. For
example, a system designed to interconnect computers over
a restricted geographical area (generally up to about one
mile) is known as local area network (LAN). One example
10 of such as system is Ethernet, which was designed by
Xerox during the late 1970's. This operates at 10Mbps and
the data are sent over twisted pairs in the form of
Ethernet frames.

For greater distances, wide area networks (WANs) are
15 employed. One form of WAN employs ATM (Asynchronous
Transfer Mode). ATM employs 53 byte cells as a basic unit
of transfer. Each ATM cell is divided into 5 bytes of ATM
layer overhead and 48 bytes of ATM payload. An ATM
network is essentially statistical in nature with the ATM
20 cells being transmitted over virtual channels established
within the network.

Through the use of a cell switching multiplexing
scheme, ATM exploits bandwidth utilization by taking
advantage of the statistical multiplexing benefits of a
25 blocking switching fabric. In a homogeneous ATM
environment, ATM terminal adapters are connected to the
ATM switching fabric, forming a potentially fully meshed
logically connected communications infrastructure. An

- 2 -

ATM adapter is any edge device which interfaces the ATM cell stream to an ATM services consumer. Typical examples of ATM adapters include ATM network interface cards (NICs) for digital computers.

5 Generally, networked computers are provided with LAN adapters for connection to a local area network, such as Ethernet. Such LAN adapters do not permit them to be connected to Wide Area Networks, such as ATM networks. While it is possible to fit special ATM adapter cards
10 into networked computers, this requires physically accessing the computers and supplying appropriate driver software.

Research has been done on ATM standardization with a view to facilitating connectivity with personal
15 computers. At the University of Melbourne, Australia, research has been carried out on running ATM cells over Ethernet. The shortcomings of this approach is that ATM is being standardized in a way that will still require two significant changes to the network, namely the
20 replacement of workstation Network Interface Cards (NICs), and modification of the NIC driver software. Bandwidth management, signalling and call management must be performed at each Sun workstation.

In order to provide a networking infrastructure with
25 maximum flexibility, there is a need to provide effective interconnectivity between computers over a wide area network while requiring minimum hardware modification.

- 3 -

According to the present invention there is provided a computer communications network comprising a network of interconnected ATM switches over which ATM cells are transmitted, and a plurality of user devices including

5 LAN interface adapters for connection to one or more local area networks (LANs), characterized in that interface means connect at least some of said respective ATM switches to said LAN interface adapters, said interface means adapting said ATM cells for transport

10 over said LANs whereby said user devices can communicate through said LAN interface adapters transparently with said network.

The communications network may, but not necessarily, be a wide area network (WAN) communicating with a local

15 area network (LAN).

In accordance with the invention the user devices, which may be personal computers, can be interconnected over the ATM network using their existing LAN adapters. The whole network, including the ATM switch fabric, thus

20 acts as a virtual LAN.

In one embodiment, the ATM cells are encapsulated in LAN frames and delivered in encapsulated form over the Ethernet LAN direct to the LAN adapter cards. In another embodiment, the interface means provide bridging,

25 network-layer forwarding and LAN emulation functions to permit transparent communication between any of said user devices over the ATM network. Such a device creates LAN

- 4 -

frames from the ATM cells and vice versa, and is known as a ridge or bridge/router.

An advantage of this arrangement is that neither the network interface adapters nor the accompanying driver
5 software at the local workstations require replacement.

The ATM network is modeled as a distributed router, which shares topology and reachability information with external routing peers. Connections within the network are set up on demand, using lightweight signaled calls
10 routed through predefined virtual paths, each containing multiple connections with similar traffic characteristics. Since each connection has only a small committed information rate, but is allowed to burst to the level of the virtual path which contains it,
15 resources are controlled while still achieving statistical gains from the aggregation of traffic.

The network of ATM switches emulates a LAN and the system functions as an extremely large, distributed bridge/router. Devices connecting to the system "believe"
20 they are connecting to a large LAN. Somewhere on the LAN there appears to be a router, through which many more networks can be reached. The devices are completely unaware of the true architecture of the system. They have no way of knowing that the LAN is being extended over the
25 ATM fabric, and that the networks behind the "router" are also part of that same ATM fabric. Routers connected to the system also view it as a LAN with an attached router. Routing information is exchanged between the external

- 5 -

router and the VIVID "router" using standard routing protocols.

The invention will now be described in more detail, by way of example only, with reference to the

5 accompanying drawings, in which:-

Figure 1 is a block diagram of a wide area network operating in accordance with a first embodiment of the invention;

Figure 2 is a block diagram of a wide area network
10 operating in accordance with a second embodiment of the invention;

Figure 3 is a diagram illustrating the function of a ridge;

Figure 4 is a block diagram showing the internal
15 operation the of the ridge in more detail;

Figure 5 shows the general traffic flow in the ridge;

Figure 6 shows the traffic flow in the Ethernet - ATM direction for the ridge;

20 Figure 7 shows the traffic flow in the ATM - Ethernet direction for the ridge;

Figure 8 shows a QMAC in more detail;

Figure 9 is a block diagram of a look-up engine;

Figure 10 is a diagram showing buffer flow in the
25 ridge;

- 6 -

Figure 11 is a block diagram of a PHY module;

Figure 12 is a more detailed block diagram of a route server for ATM-LAN network;

Figure 13 shows a prior art OAM processing resource
5 for an ATM switch;

Figure 14 shows an OAM processing resource according to one embodiment of the invention;

Figure 15 shows an arrangement for connecting LAN adapters to an ATM network without offering the full
10 benefits of ATM;

Figure 16 illustrates a typical example of an interconnected system of ethernet-attached computers and ATM-attached computers;

Figure 17 schematically illustrates the relevant
15 internal features of the layers described in Ethernet-attached end stations using ATM services; and

Figure 18 schematically illustrates the highlights of a virtual ATM switch shown in Figure 16.

Referring now to Figure 1, ATM switches 1a, 1b, 1c
20 define a wide area network Asynchronous Transfer Mode network (WAN). Switch 1a is connected through network interface card (NIC) 2 to route server 4, and through network interface card (NIC) 3 to system manager 5.

Workstation 12 is connected in a conventional manner
25 via Network Interface Card (NIC) 13 to the Wide Area

- 7 -

Network. NIC 13 is specifically adapted to connect the workstation 12 to the ATM network.

Switches 1a, 1b, 1c are also connected through WAN-LAN interface devices 6, known as ridges, to router 9, workstation 8, hub 10, bridge 11 and SNMP manager 12, which each have Local Area Network adapters for connection to a Local Area Network, in this case Ethernet.

Devices 8, 9, 10 and 11 are connected to respective Ethernet ports of ridge 6a, and SNMP manager 14 is shown connected to one of the Ethernet ports of ridge 6b.

The ATM wide area network operates under control of the route server 4, which translates from the connection-less model of traditional LANs to the connection-oriented model of ATM. It implements traditional routing-table computation protocols (e.g., RIP and OSPF) and communicates with external routers to learn the extended topology of the connection-less network. It also discovers the presence of all other external devices to complete its knowledge of the topology of the connection-less network.

The route server 4 learns the topology of the ATM network from the System Manager 5 and configures the ridges 6, as required to map between the ATM topology and the LAN topology.

The function of the route server 4 is to maintain the topology of the devices attached to the ATM-LAN

SUBSTITUTE SHEET

- 8 -

system. The topology information is used to forward and route LAN traffic from source to destination over an ATM network. Traffic forwarding is done by the route server and also by the bridges 6 using information that the route server conveys to them about the topology. This ability to allow devices to connect anywhere in the network that the configuration allows is unique to this system, and the function is performed by the route server topology management.

10 The route server 4 thus contains the administrative information that defines a virtual LAN, such as the LAN network numbers and device membership, filters and access restrictions.

15 The route server knows the layout of switches and trunks, and discovers the presence of all attached devices. The route server uses this complete network knowledge to ensure that packets are forwarded through the system correctly, using a dynamic mesh of ATM connections.

20 In smaller networks, the route server 4 also acts as an ARP Server, and a Broadcast Server (larger networks will contain several separate Route, ARP, and Broadcast Servers). In this role, the route server 4 forwards broadcast traffic to any other network elements that need to receive it. This allows the route server 4 to employ various heuristics and optimizations to limit the amount of broadcast traffic flowing through the network. For example, ARP requests can often be handled directly by

- 9 -

the Route Server, which already knows the addresses of most devices present in the network. For those packets that need to be flooded, the Route Server can send a single copy of the packet to each Ridge that the
5 broadcast is intended for, along with a mask to indicate the ports over which the packet is to be flooded.

A simplified block diagram of the route server 4 is shown in Figure 12. The route server 4 comprises a Centralized Routing and Forwarding Server 400, a NIC
10 device driver 407, a transaction manager 402, a topology manager 403, a multicast server, 404, a route manager 405, and an ADP manager 406. Devices 402 to 406 are connected to forwarder 400 and SNMP agent 408. Transaction manager 402 is connected to standby server
15 409. Forwarder 400 is connected directly to SNMP agent 408.

The route server 4 maintains intelligence about the location of all devices on the system. This information is learned dynamically as devices attach to the network
20 and is learned through communication via standard Routing Protocols, e.g. IP and IPX with Routers on the edge of the system. The intelligence is distributed to ridges 6 via an Address Distribution Protocol. The communication with the ridges 6 is allows the ridges to perform data
25 forwarding directly on the majority of data.

The Centralized Multicast Server 404 processes all multicast traffic. Where possible, the route server 4

- 10 -

responds to the sent data without further broadcasting into the network.

The route server 4 also carries out LAN topology management to dynamically allow Adds, Moves and Changes of LAN devices; and to perform Admission Control on devices against rules configured by the System Administrator. The devices 9 to 14 attached to an ATM LAN are dynamically discovered by the route server and admitted into the network if allowed. The ridges 6 participate in the discovery by detecting the device. The route server determines if admission is allowed and if so what services are allowed. The route server 4 also maintains knowledge about the location of the device for long periods of time. The route server 4 allows devices of dissimilar networks to share ports into the system.

The route server 4 also provides Flexible, portable and redundant platform support. The route server is run on a SUN workstation with an ATM network interface card supporting both single and multiprocessor platforms. A redundant route server 409 is supported and takes over in case of failure of the primary route server 400. An ATM-based messaging protocol is used between the two platforms to ensure coordination.

Returning to Figure 1, each ATM Switch 1 provides high-bandwidth cell-switching that is the core of the System. There are three types of ATM switch, namely an ATM Workgroup Switch (WGS), a 36170 Switching Shelf and a 36170 Access Shelf.

SUBSTITUTE SHEET

- 11 -

The WGS is a low-cost, 12-port, 1.6 Gbit/s ATM switch for use in Customer-Premise applications. The 36170 Switching Shelf is a 12.8 Gbit/s ATM switch which interconnects up to 8 ATM Feeder Switches and/or Access
5 Shelves. The 36170 Access Shelf is a 12-slot, 1.6 Gbit/s ATM switch. Each slot can contain one of many ATM interfaces. The WGS is described in more detail with reference to Figures 13 and 14.

Currently, ATM switches that provide OAM (Operation
10 and Management) support at an ingress port require a dedicated microprocessor 52 to interpret and generate OAM cells. OAM cells are transferred to/from the microprocessor 52 by cell processing engine 50 using shared-RAM 51 as shown in Figure 13.

15 The drawback to this approach is that additional cost and complexity (PCB area, addition components, shared RAM systems) is required to support the OAM functionality. This addition cost hinders the ability to make a cost competitive, full featured, multi-port
20 switch.

In accordance with the invention as shown in Figure 14, the cell processing engine is modified to redirect all relevant OAM cells to a centralized OAM processing resource using a preconfigured internal switch address.
25 All processing of the cells occurs at his one microprocessor, and this eliminates the need for a dedicated microprocessor 51 on each port, as shown in Figure 14.

- 12 -

In one suitable scheme, cells with VCI (Virtual Channel Identifier) = 3 or 4 (segment and end-to-end) and VP (Virtual Path) switched compose the F4 (VPC (Virtual Path Connection)) OAM Flow. Cells with PTI (Payload Type Identifier) = 4 or 5 (segment and end-to-end) and VC switched compose the F5 (VCC (Virtual Channel Connection)) OAM Flow.

Ingress F4/F5 OAM Flow cells with an OAM Cell Type = 0001 (Fault Management) and a Function Type = 0000,0001 or 1000 (AIS, FERF, Loopback) are extracted from the cell stream and redirected using a predefined Header. These cells are the one that are processed by the centralized OAM processor.

A suitable centralized OAM processing resource is a CCM (Control Complex Module), which is responsible for providing OAM support to all UCS (Universal Card Slot) modules. The UCS modules direct the required ingress OAM cells [VC AIS (Virtual Channel Alarm Indication Signal), VC FERF (Virtual Channel Far End Receiver Failure), VP AIS (Virtual Path Alarm Indication Signal), VP FERF (Virtual Path Far End Receiver Failure), Segment Loopback, End-to-End Loopback] to the CCM. The CCM generates the required egress OAM cells.

Prior art ATM switches that implement UPC (Usage parameter Control), ATM address translation or other enhanced functionality at high-speed (155 mbps or greater) ingress or egress ports require a dedicated microprocessor to update port specific parameters located

- 13 -

in a dual-port/shared memory system. These parameters may include new VPI/VCI's, proprietary switch headers, UPC parameters, statistics and connection validity. When a cell arrives at the ingress or egress port, a hardware
5 (H/W) engine then uses these parameters to process the cell appropriately.

The drawback to this approach is that additional cost and complexity (PCB area, addition components, shared RAM systems) is required to implement the enhanced
10 functionality. This additional cost hinders the ability to make a cost competitive, full featured, multi-port switch.

Enhanced functionality is added to the H/W cell processing engine to eliminate the need for the
15 additional processor and shared memory system. To do this, the data and control streams are amalgamated and the H/W cell processing engine is designed to interpret the control cells. This updates the RAM as required to configure ingress and egress connections, and sends
20 responses to status requests. This can be achieved providing a Field Programmable Gate array, such as one available from Xilinx corporation, in the cell processing engine. The cell processing engine then can interpret the control cells.

25 An additional benefit of this approach is a reduction in the bandwidth requirement of the RAM system. Existing implementations require high-speed SRAM to implement the shared memory system. By eliminating the

- 14 -

requirement to provide additional accesses by the external microprocessor, the RAM bandwidth may be reduced significantly.

This technology may be used by any ATM switch. In addition, this approach can be used by any line card in a switch environment requiring frequent parameter updates.

Returning now to Figure 1, the ATM switches 1a, 1b, 1c are connected to the respective ridges 6a, 6b, 6c. In one embodiment, each ridge 6 has twelve Ethernet ports 7 for connection to an Ethernet LAN, Ethernet adapters of a local workstation 8, a conventional router 9, Hub 10 or bridge 11.

The ridges 6 carry out the bridging, network-layer forwarding and LAN emulation functions to permit transparent communication between any of the user devices over the ATM network. The ridges permit the devices designed to communicate in the LAN environment via the ATM environment.

In fact the ridges carry out six basic steps as follows:

(1) Source Address Validation

When a packet arrives on a ridge LAN port, the ridge verifies that it has seen the source MAC address on the port before. If the address is found in the source address table for the port, the packet proceeds to Destination Identification.

- 15 -

If the address is not found, a new station has entered the system, and the packet is forwarded to the Route Server for processing.

(2) Destination Identification

5 Once the Ridge has validated the source address of the packet, it examines the destination MAC address of the packet. Several possibilities can occur at this point:

- The packet is addressed to a broadcast address,
10 or to a MAC address that is not in the Ridge's destination address table. The packet is forwarded to the Route Server for processing.

- The packet is addressed to the MAC address of the VIVID "router" itself. This implies that the packet
15 is to be forwarded, so the network layer address of the packet is examined. If the network layer address is in the destination address table, forwarding information is retrieved from the table, and the packet proceeds to the Filtering stage. Otherwise, the packet is sent to the
20 Route Server.

- The packet is addressed to a MAC address that is in the Ridge's destination address table. In this case, the packet can be bridged, so the forwarding information is retrieved from the table, and the packet
25 proceeds to the Filtering stage.

The purpose of this step is to map the logical destination address in the lookup table to a physical

- 16 -

address for the output media. The forwarding information retrieved from the destination address table therefore depends upon how and to where the packet is being forwarded. If the destination is attached to a port on the Ridge, and the packet is being MAC layer forwarded, the destination port ID is retrieved from the table. If the destination is attached to a port on the Ridge, and the packet is being network-layer forwarded, the destination port ID, and the MAC address of the destination are retrieved. Finally, if the destination is located on another Ridge, only the ATM address of the egress Ridge is retrieved from the table.

The physical-to-logical connectivity mapping is performed at this point by associating a geographically based physical layer ATM address with the logical network layer destination address within the packet.

(3) Filtering

After the source and destination addresses of the packet are verified, the Ridge checks to see if the devices are allowed to talk to each other. Generally, this is simply a check of source address, destination address, and protocol type, although it may be necessary to look deeper into the packet to perform application-level filtering when required. If the devices are not allowed to communicate, the packet is dropped. Otherwise, the packet proceeds to the Transformation stage.

(4) Transformation

- 17 -

If the packet entered the system through a LAN port, and is being network layer forwarded, the Ethernet, or 802.3 encapsulation is removed, along with the source and destination MAC addresses in the packet, and the time-to-
5 live for the packet is decremented. The packet is then encapsulated as a routed pdu, to travel over the ATM fabric.

When the packet is forwarded on a Ridge port, the source MAC address of the packet is set to the MAC
10 address of the VIVID "router". The destination MAC address is set to the value that was retrieved from the destination address table during Destination Identification. Then, the packet is encapsulated in the appropriate format for the LAN.

15 (5) Call Setup

If the packet is to be forwarded over the ATM fabric, the Ridge checks to see if it has an SVC to the egress Ridge. If an SVC has not already been established, a call is originated to the egress Ridge, using the ATM
20 address that was obtained during Destination Identification.

(6) Transmission

The packet is transmitted on the Ridge egress port, or fragmented into ATM cells and sent out over an SVC to
25 the egress ridge.

Each ridge 6 therefore provides media-rate interconnection between a traditional Local Area Network

- 18 -

(LAN) (e.g., Ethernet/802.3, Token Ring/802.5, Fiber Distributed Data Interface (FDDI), etc.) and the Asynchronous Transfer Mode (ATM) wide area network.

The ridges 6 may be considered a 178,560 pps switching and concentration element. The functions performed by the ridge are, more specifically, :

- a) ATM layer segmentation and re-assembly,
- b) ATM Adaptation Layer 5 (AAL5),
- c) encapsulation and de-encapsulation of frames in RFC 1483 headers,
- d) derivation of an ATM address and/or VPI/VCI from the header of a frame,
- e) Ethernet bridging or network-layer forwarding of frames,
- f) fragmentation of IP frames received on the ATM port, and
- g) Transmission Convergence sublayer processing compliant with SONET STS-3c.

As will be described in more detail below, on the Ethernet interface the ridge encapsulates the received cells and sends them into the ATM network toward their destinations. On the ATM interface, the ridge encapsulates cells destined for its attached Ethernet hosts before sending them over the Ethernet to their destinations.

- 19 -

To achieve fairness without contention, a token is continuously passed between the ridge and its attached Ethernet hosts. When a host receives a token, it sends a frame, or a set of frames, to the ridge. Each frame
5 contains one or more ATM cells. After sending a predetermined number of frames, the host passes the token back to the ridge, which may send a set of frames itself, before returning the token to a host. In this way the ridge and host(s) are synchronized so that each host gets
10 a specified amount of bandwidth without contention or collision. By varying the number of frames each host sends before releasing the token, or varying the token passing protocol itself, each host can receive a different grade of service.

15 Figure 2 shows an alternative arrangement, in which like parts are referenced with like reference numerals. In Figure 2, ATM switch 1c, as well as being connected to ridge 6c, is connected to ridge-like gateway 14, which in turn is connected to workstation 15. Gateway 14, instead
20 of converting ATM cells to Ethernet format, and vice versa, encapsulates the incoming ATM cells in Ethernet frames, thus allowing them to be received directly by the Ethernet adapters in the local workstations with the aid of Ethernet drivers.

25 Referring now to Figure 3, the ridge 6 illustrated has twelve 10 Mbps (10baseT) Ethernet ports 20 for connection to Ethernet devices, a single RS-232 serial port 21, and a 155Mbps one OC-3 over Multi-Mode Fiber ATM

- 20 -

port interface port 22. As indicated above the ridge 6 performs bridging, LAN emulation and network layer forwarding functions. Both user data and control traffic (to and from the route server and System Manager) are
5 carried on the ATM interface.

The function of the ridge 6 is to perform Ethernet bridging, network-layer forwarding and LAN emulation for 12 Ethernet ports and a single ATM port. Traffic bound between Ethernets may be either bridged or network-layer
10 forwarded directly; when the ATM port is required for traffic to other ridges, encapsulation and ATM-layer processing is performed.

Referring now to Figure 4, which is a more detailed block diagram, the ridge 6 comprises Ethernet MAC
15 receiver 23 and transmitter 24, outgoing frame processor 25, incoming frame processor 26, incoming and outgoing frame storage memories 27, 28, look-up engine filter 29 (illustrated in Figure 9 and described more fully in our co-pending application no. PCT/CA94/00695, filed December
20 22 1994 and entitled Look-up Engine for Packet-Based Network), ATM layer segmenting processor 30, cell autopadder 31, TC layer SONET framer 32, ATM layer re-assembling processor 33, and fast queue servicing controller 34. TC layer SONET framer is connected to ATM
25 module 35 providing the ATM interface 22 for connection over optical fiber or twisted pair copper wire.

The purpose of the fast queue servicing controller 34 is to enqueue on an output queue packets destined for

- 21 -

Ethernet ports. There is an output queue for each port. When a queue is serviced, the packet is transferred to the Ethernet MAC port, which must be ready to accept the packet. All queues have equal priority.

5 In the fast queue servicing controller 34, the availability of ports and data is factored into the request for the next port. Port availability is presented as a bit mask in the address. A priority encoded with round robin priority ensures that if the current queue
10 (as defined by the round robin) cannot be serviced (because either the port or data is not available) then the next highest priority port with all the requirements satisfied is returned. This device always therefore returns a useful result. With the address mapped bit map
15 of useful data, the result can be returned with just one read operation, thereby improving performance.

 The fast queue servicing controller 34 can in fact be applied to any set of parallel queues that require servicing, and it is not limited to a round-robin
20 servicing mechanism, but can also be implemented where unequal priorities are associated with the queues.

 The traffic flow in a ridge 6 is shown in Figure 5. The ridge 6 is designed for separate data paths in the receive and transmit directions. The only exception is
25 locally switched Ethernet traffic, which is routed through the Segmenter RAM system via bypass unit 45. This approach eases the requirement for a single memory system which would quickly become the system bottleneck.

- 22 -

Quad MAC 40 providing transmitter 24 and receiver 23 in Figure 4 has 3 Kb input and output FIFOs 45, 46, connected to segmenter RAM 41 and segmenter 42 in the ATM direction side and re-assembler RAM 43 and re-assembler 44 on the Ethernet side input. Bypass unit 45 allows incoming Ethernet traffic to bypass the ATM network and pass directly to the output of Quad MAC 40.

Figure 6 shows in more detail the packet flow in the ATM direction. First the Ethernet frames are buffered in their entirety inside the Quad MAC ASIC FIFO 45, then the QMAC 40 notifies the AXE RISC processor 48 via Reception Controller 47 that a DMA is required. The AXE (Transfer Engine) 48 initiates DMA to Segmenter RAM without taking into consideration which port is selected, and the Reception Controller 47 selects the port using a round-robin priority scheme.

The Look-up Engine 29, referred to above and described in our co-pending application then loads the frame header information in "fly by" mode and starts context searching for source, destination MAC address, protocol type, port group, etc.

The AXE 48 takes the look-up results, and may drop the packet or perform a network-layer transformation if instructed. It then reformats the packet as required into a CS-PDU and notifies the Segmenter 42 to start cell slicing. Segmenter performs ATM segmentation and returns the buffer when complete.

- 23 -

The AXE 48 is a 50 MHz R3000 RISC engine, with an integral high speed DMA and a separate, secondary processor bus. In order to maintain media speed performance, it must complete the tasks described below in 5.6 us (280 cycles). The AXE 48 keeps the DMA data flowing between the QMACs and Segmenter RAM in the foreground, moving about 512 bytes at a time; there is no interleaving of ports once a packet DMA has begun. In the background, packet information is retrieved from the LUE FIFO, and AXE packet processing begins.

The packet is either discarded, bridged or network layer forwarded. For bridged packets no packet modification is performed.

For network layer forwarded packets, fields in the network layer payload header of the packet are modified; in the IP case, for example, the TTL is decremented and the checksum adjusted. The new destination MAC address is appended to the packet descriptor for insertion by the Transmit Controller. The processing at this stage is protocol dependent.

The packet is encapsulated in RFC 1483 type LLC encapsulation for ATM Adaptation Layer 5

The VC is inserted from the LUE 29.

The frame is queued to the Segmenter for ATM or local transmission.

Ridge Host Processor originated traffic simply appears as a "13th port" to the rest of the ridge (Figure

SUBSTITUTE SHEET

- 24 -

8). Part of the Host Processor complex includes a FIFO interface which mimics the Quad MAC 40.

Data flows out of the Ethernet ports in a symmetrical path to that of the previous description as shown in Figure 7, which shows data flow in the ATM - Ethernet direction. This is similar to that described in Figure 6 except that transmit control processor 50 conducts QMAC Direct Memory Access and adds MAC headers. The Transmit Controller 50 is functionally similar to the AXE 48, as packets are dequeued from the Segmenter and Reassembler RAMs, MAC addresses are inserted, and they are DMA'd to the Quad MACs. Like the AXE, the TXC is implemented with a 50 MHZ R3000-based RISC processor, and must complete its processing in 5.6 μ s.

There are three sources of Ethernet packets transmitted from the ridge: ATM cells, local inter-port traffic, and local Host-processor originated. In the latter two instances, the frames are in Segmenter memory rather than Reassembler memory. As such, the Transmit Controller 50 must deal with two traffic sources to feed the QMACs 40.

The packet flow from the Reassembler is as follows:

- 1) The ATM cells are reassembled into Reassembler RAM;
- 2) The Reassembler DMAs the packet header into the LUE, which determines the source and destination MAC addresses, as well as the outgoing portset. The information is appended to the frame for use by the

- 25 -

Transmit Controller;

3) The frame is queued to the Transmit Controller by the Reassembler.

The packet flow from the Segmenter memory system simply involves the Segmenter queuing frames to the Transmit Controller 50. After the Transmit Controller RISC processor is notified of the arrival of frames via the ATM or local switching paths, it inserts the MAC addresses and conducts the DMA out to the appropriate port of the Quad MAC.

The QMAC 40 is shown in more detail in Figure 8. This comprises an ASIC which streams four ports of Ethernet traffic to and from high speed, 32 bit wide synchronous data buses; an external DMA is required. The QMAC 40 features integral 10BaseT or AUI transceivers, a full compliment of packet and byte counters, and an internal 3k packet buffer in each direction

The look-up engine 29 is shown in Figure 9 and described in our co-pending application referred to above. The lookup engine (LUE) is used each time a packet is received off the Ethernet or the ATM network. The type of information that the engine provides, which is held in a look-up table, depends on the direction of packet flow and the type of packet. The LUE will provide all the information needed to find the path to each known destination, as well as provide default information in the case of unknown destinations. The default information

- 26 -

in most cases will direct the packet towards the route server.

The LUE 29 is based on table lookups using nibble indexing on variable portions of the packet, such as MAC and network layer addresses, and bit pattern recognition on fixed portions for network layer protocol determination. Each lookup table is organized into a hexadecimal search tree. Each search tree begins with a 16 word root table. The search key (e.g. MAC address) is divided into nibbles which are used as indices to subsequent tables. The 16 bit entry in the table is concatenated with the next 4 bit nibble to form the 20 bit address of the next 16 word table. The final leaf entries point to the desired information. A block diagram of the LUE is shown below:

Bit pattern recognition is achieved by a microcode instruction set. The microcoded engine has the ability to compare fields in a packet to preprogrammed constants and perform branches and index increments in a single instruction typically. The microcode engine has complete control over the search algorithm, so it can be tailored to specific lookup functions, and microcode is downloaded as new functions are required. The output of the microcoded engine packet parsing is an index which the AXE can use to rapidly vector to a processing routine.

Learning and aging of source addresses requires tree manipulation, and is accomplished as a background task

- 27 -

with the help of the AXE and Host Processor. Newly discovered source MAC frames result in an internal message forwarded to the Host Processor, requesting parameters be added to the LUE Source RAM. The Host
5 Processor then accesses the LUE RAM through a dual-porting mechanism and re-arranges the tree as required.

The LUE is physically partitioned into a large FPGA, separate source and destination lookup memories of 512k kB and 1 MB respectively, a downloadable microcode RAM,
10 and a combination of Xilinx and FIFO devices to interface to the RISC processors.

To ease the design requirements of a single, superfast memory system required for Segmentation, Reassembly, and QMAC traffic, the ATM SAR function is
15 split in two by the ridge. As such, the architecture is horizontally divided by direction, and is truly a full duplex system.

The Segmenter complex consists of an ATMizer, an interface into the Segmenter RAM on its primary bus, some
20 additional high speed pointer memory on its secondary bus, and a dual-ported memory to the Reassembler. The latter is used as a SAR communication buffer for OAM and lightweight switching support.

A dedicated ATMizer is used to implement ATM
25 reassembly in the ridge of up to 1024 simultaneous VCs. As packets are reassembled, the LUE snoops the incoming first cell, and provides the destination MAC address if it's a routed PDU (the source MAC is automatically the

SUBSTITUTE SHEET

- 28 -

route server). The MAC is appended to the packet by the Reassembler for quick insertion by the Transmit Controller.

Cells may be passed from the Reassembler to the Segmenter using an inter-ATMizer cell link. This is useful for loopback diagnostics.

The Segmenter and Reassembler RAMs are similar in function, and provide the intermediate CS-PDU buffering between the ATM and LAN environments. The Segmenter memory has 4 ports: the AXE Secondary bus, the Quad MAC (AXE Primary bus), the Transmit Controller and Segmenter. The Reassembler memory has 3 ports: the Transmit Controller, and both the Primary and Secondary buses of the Reassembler. Both are 512 kB in size.

Each RAM system has a similar arbitration controller, and services access requests one at a time in a round-robin fashion. The arbiter allows addresses to accumulate behind buffers such that switch-over time between interfaces is minimized. Currently the memory systems use high speed 15 ns RAMs to achieve 3 cycles for arbitration and 4-5 cycles for write and reads, respectively. The RAM systems provide over 800 Mbps in sustained bandwidth.

The ridge requires a supervisory Host processor to handle system startup, power-on diagnostics; downloading of LUE and RISC processors; run Spanning Tree algorithm; perform network management (SNMP, 4602 I/F); perform local serial port configuration; connection management

- 29 -

(Q.2931); perform "learning" and table maintenance for the LUE; run PHY module state code if necessary (OC3); and control faceplate LEDs

The Host processor complex is realized with a 6 MIP
5 MC68349 CPU operating at 25 MHz, with the following
peripherals: 2 MB RAM; two banks of 1 Mbyte Flash EPROM
(Intel based); 32 KB Battery backed NVM (Non-volatile
Memory); 2KB FIFOs in each direction for packet
transmission; ATMizer serial downloader; and RS-232
10 transceiver for local serial port configuration.

A 5-processor pipelined architecture like the ridge
can be expected to have significant data-passing
requirements; there are essentially three, as described
in the following sections.

15 The Host Processor needs to pass configuration,
maintenance, and statistics polling packet messages to
each of the ATMizers, as well as communicate with
external Ethernet-based devices. To streamline these two
requirements, the Host Processor appears as a "13th port"
20 in line with the QMACs to the AXE and Transmit
Controller. Frames destined for internal processors are
flagged and loaded into the Segmenter memory by the AXE,
in much the same way as regular traffic. To get to the
Reassembler, a message packet must be physically recopied
25 by the Transmit Controller into Reassembler memory.

In the other direction, the Transmit Controller
differentiates between local messaging and MAC-addressed
traffic destined for the Host by prepending a flag to the

SUBSTITUTE SHEET

- 30 -

packet. The Host processor uses its integral DMA controller to move data to/from its QMAC-emulation FIFOs.

As buffer pointers are passed between the ridge ATMizers, a simple interrupt flag is used to signal their requirement for service. At the same time, the "CPCond" input is set, such that a single R3000 branch instruction can efficiently poll and vector on the flag in a single cycle.

The Segmenter receives an interrupt from the AXE indicating the arrival of a data buffer

The TXC receives an interrupt from the Segmenter and Reassembler indicating the arrival of a data buffer

The Segmenter and Reassembler interrupt each other to service their dual-ported memory queues

The Segmenter and Reassembler need a direct communication path for bi-directional protocols such as lightweight switching and ATM OAM. This is provided by a small dual-ported memory between them, and the capability to interrupt each other.

With four processors moving data around the ridge, an optimal scheme for passing buffer pointers and returning free ones is essential. With the multiporting of Segmenter and Reassembler memories, no actual copying of data is necessary. For simplicity and low cost, buffer pointers are passed through the multiported memories in the same path as the data. To avoid complexity in the

SUBSTITUTE SHEET

- 31 -

returning of buffers, they are passed in a unidirectional manner, as shown in Figure 10.

Packets received by the AXE are forwarded to the Segmenter, whether they are destined for the ATM network or not. In the case of local traffic, the Segmenter will requeue the packet to the Transmit Controller, which means that the returned free buffer list is managed by a single processor. This avoids re-entrance problems and simplifies overall buffer management.

10 Data buffers received by the Reassembler can simply be passed to the Transmit Controller for Ethernet transmission and returned when the DMA is complete.

Statistics are collected by a combination of hardware and software in the ridge, depending on the application. The Quad MACs have hardware counters for:

- Number of frames sent/received
- Number of bytes sent/received
- Ethernet reception errors (CRC, alignment, runt)
- Number of collisions
- 20 • Number of failed transmissions due to >16 collisions
- Other statistics are gathered by ATMizers for reporting via the Host Processor. Examples of these are:
- 25 • Number of cells sent/received

SUBSTITUTE SHEET

- 32 -

- Number of security violations/port

The ridge PHY module addresses the current debate over physical interface types by making it easily changeable. This is shown in more detail in Figure 11.

- 5 The ridges are therefore a key element of the system, concentrating 12 Ethernet ports with local switching onto a single ATM trunk for connection to an ATM switch. The particular embodiment of the ridge described features: 12 10BaseT Ethernet ports; one
- 10 Modular ATM port (first release: 155 Mbps STS-3c over Multimode fiber optic cable); Full media-speed bridging and network-layer forwarding between all Ethernet and ATM ports; Broad range of packet filtering capabilities; Supports IP fragmentation in the ATM to Ethernet
- 15 direction; Manageable using SNMP and standard MIBs; Spanning Tree algorithm; ATM support: AAL5, CLP, OAM, traffic shaping, Newbridge "lightweight switching"; Remote traffic monitoring with Ethernet "snooping" mode; Rack, desktop, or wall mountable in standalone packaging;
- 20 Local management interface (serial port); Software downloadable using "Flash" memory; "Soft" RISC and ASIC-based design: over 200 MIPS

- The described ridge is thus an ATM LAN emulator designed to marry a routed Ethernet environment to an ATM
- 25 WAN network with the following features: Low cost; Full media speed- bridging and network layer forwarding; Accommodate standards yet to be defined/invented; Support a full range of filtering capabilities; Understand and

SUBSTITUTE SHEET

- 33 -

take advantage of routed protocol capabilities; Work standalone, or integral to a variety of ATM hubs; Straightforward migration path to ASICs.

Key elements of the ridge are:

5 High density, fast Ethernet MACs.

Powerful table searching engine. A fundamental aspect of bridges and routers is the ability to quickly associate ports, VCs, MAC addresses, security restrictions, etc., in real time. On the ridge, this
10 necessitates a lookup every 5.6 us into a table holding at least 8,000 MAC entries.

Flexible network layer hardware parser. To perform high speed network layer forwarding, frames need to be rapidly parsed and modified with new MAC addresses, TTL
15 and checksums. A high speed RISC processor actually does the work and is coupled with the table-searching engine to form a frame-processing complex.

Fast, multiported ATM to Ethernet buffering system. As the ridge is a store-and-forward device, and
20 congestion is possible in both directions, the amount of internal buffer storage is a consideration. The memory system must be fast, providing at least 275 Mbps in each direction, relatively large, and low cost. Other
important characteristics include efficient arbitration
25 of the RAM systems and low access latency.

A fundamental aspect of the ridge is the repackaging of Ethernet frames, insertion of a VCI, and dissection

SUBSTITUTE SHEET

- 34 -

into 53 byte cells. The reverse process requires simultaneous piecing together of frames, and redirection to the appropriate Ethernet port. The cell size, and ATM processing required, is highly dependent on the target
5 environment.

The "ATMizer" has a 50 MHz R3000 core coupled with internal RAM, caches, DMA and ATM serialization hardware required to process ATM cells in real time.

RISC-based DMA. The ability to move ridge data
10 around at high speed, modifying fields and managing queues, calls for an extremely fast and cost-effective RISC solution. To streamline the development process, and take advantage of a highly optimized integrated solution, the ridge re-uses 2 ATMizers in a non-ATM application to
15 perform rapid data manipulation.

High speed ATM port. Connection to an ATM switch via transmission medium (standalone ridge) or a suitable backplane interface (Stealth in the 36150 switch). As the choice of physical medium and framing scheme is still
20 being actively debated, a swappable PHY module is used for this function.

A relatively low speed processor is required as a Housekeeping Processor for network management, diagnostics, configuration, and overall supervision. In
25 one embodiment, this is a Motorola 68349 microprocessor clocked at 25MHz.

SUBSTITUTE SHEET

- 35 -

A further aspect of the invention will be described with reference to Figures 15 to 18.

In the Figure 15, the ATM attached digital computer and legacy LAN adapters 100, 101 are used to interface
5 legacy Ethernet LANs and digital computers to the ATM switching fabric 102, which may be composed of one or more ATM switches 103. This configuration is used to transport LAN data traffic between the various termination units. However, the benefits of ATM, which
10 include guaranteed bandwidth (BW), traffic management, and so on, are not extended past the legacy LAN adapters to the legacy LAN attached (e.g., Ethernet) digital computers. Thus, it is impossible to use Ethernet, in the traditional carrier sense multiple access with
15 collision detection (CSMA/CD) mode of operation, to provide all the services available to directly attached ATM components (e.g., digital computers), unless a meta level access control mechanism is employed in the legacy LAN environment.

20 Providing ATM services to an Ethernet-attached end station (digital computer), which is interfaced to the ATM switching fabric via an ATM-attached legacy LAN adapter requires that a connection management, signalling, and bandwidth management mechanisms be
25 provided, which deterministically control access to the Ethernet LAN from all end stations attached to it, and extends ATM Q.2931 signalling to the Ethernet end stations.

SUBSTITUTE SHEET

- 36 -

The provisioning of full ATM services to Ethernet-attached terminal devices, using digital computers as a specific example of said terminal devices, will now be described.

5 Figure 16 illustrates a typical example of an interconnected system of Ethernet-attached digital computers 101 and ATM-attached digital computers 100.

The Ethernet-attached consumers of ATM services are depicted as blocks containing four layers:

10 "Application(s)", "ATM API", "802.3 Driver", and "802.3 Interface", respectively 101a, 101b, 101c, 101d, each attached to Ethernet backbone 104, which is attached to ATM switch 103 by a virtual ATM switch 105 to be described in more detail below.

15 The "Application(s)" layer 101a is intended to schematically represent an application, which is layered upon an ATM application program interface (API), which is used to extend ATM services to the application.

20 The "ATM API" layer 101b is intended to schematically represent a component providing API services to the upper layer consumer, in terms of the lower layer provider, which is the "802.3 Driver" in this figure. Typical services provided by this API includes "atmConnectReq", which is used to request that a
25 connection be established to an ATM endpoint corresponding to an application specified ATM address; "atmTx", which is used to transmit service data units (SDUs) to an ATM end point, over a previously established

SUBSTITUTE SHEET

- 37 -

connection; and "atmRx", which is used to receive SDUs from an ATM end point, over a previously established connection. The top portion of Figure 17 schematically represents these constructs.

5 The "802.3 Driver" 101c layer is intended to schematically represent a component which converts SDU's into ATM cell format, and provides the ATM connection management, bandwidth management, and signalling. This layer is used to create and destroy ATM connections,
10 convert SDUs into ATM cells for transmission, and ATM cells to SDUs for reception, and controls the transmission of ATM cells onto the Ethernet LAN.

 The "802.3 Interface" layer is intended to schematically represent the ethernet interface to the
15 ethernet LAN. This interface is used to transmit and receive Ethernet frames.

 The ATM-attached consumers of ATM 100 services are similarly depicted as blocks containing four layers: "Application(s)", "ATM API", "ATM Driver", and "ATM
20 Interface", respectively 100a, 100b, 100c, 100d.

 The upper layer interface to the "ATM API" 100a provides the same interface to an application residing on an ATM-attached end station as the "ATM API" provides to the ethernet-attached end station deriving ATM services.

25 The "ATM Driver" layer 100b is intended to schematically represent a component which converts SDU's into ATM cell format, and provides the ATM connection

SUBSTITUTE SHEET

- 38 -

management, bandwidth management, and signalling. This layer is used to create and destroy ATM connections, convert SDUs into ATM cells for transmission, and ATM cells to SDUs for reception, and controls the
5 transmission of ATM cells into the ATM network.

The "ATM Interface" layer 100c is intended to schematically represent the ATM interface to an ATM network. This interface is used to transmit and receive ATM cells.

10 In Figure 16, an ATM switch is represented by an icon resembling an X. Two instances of this icon are present, 103 and 103a. Switch 103 provides direct links to ATM end points. Switch 103a forms part of virtual switch 105, and is layered on top of an "802.3 Driver"
15 layer.

In Figure 16, there is an implied peer relationship between all the instances of the ATM application(s). In other words, the applications are capable of exchanging information on a peer basis, by utilizing the services
20 made available via the ATM API.

ATM services are extended to the Ethernet-attached end stations, using the ethernet LAN as a virtual port extension to the virtual ATM switch 105. An ethernet host wishing to exchange ATM messages with any other end
25 station connected in Figure 16 forwards the cells to be transmitted to the ATM virtual switch, which in turn forwards the cells to the appropriate ATM link. It is important to note that even if the ATM end station to

SUBSTITUTE SHEET

- 39 -

which the cells are to be transmitted to is on the same Ethernet segment as the originator of the cells, the cells are still forwarded to the virtual ATM switch first.

- 5 The lower portion 105b of the virtual switch 105 can be provided by a ridge as described above.

 In order to ensure deterministic access to the ethernet LAN, any station other than the virtual ATM switch, cannot transmit any data on the LAN segment until
10 it receives a management indication frame from the virtual ATM switch. The management indication frame contains information which specifies which virtual channel connections (VCCs) the Ethernet-attached end station can transmit on. Also, given that the virtual
15 ATM switch is responsible for generating the management indication frames, it meters data to be transmitted from the virtual ATM switch to the Ethernet-attached end stations according to the bandwidth parameters associated with the VCC, when it was initialized.

- 20 Figure 17 schematically represents relevant internal features of the layers described in the Ethernet-attached end stations using ATM services. The right hand side of the figure delineates the layers for reference purposes. An application developed to utilize the ATM services in
25 the Ethernet-attached end station would use the interface points depicted at the top of the ATM API. The ATM API is not exhaustive in terms of the service interfaces shown and previously referred to. A representative

SUBSTITUTE SHEET

- 40 -

subset are shown. Figure 17 illustrates the data flow from SDUs to ATM cells to Ethernet frames. The data flow of management frames (indications) is also illustrated.

Figure 18 schematically illustrates the highlights of the virtual ATM switch 105 referred to in Figure 16. Figure 18 illustrates the centralized bandwidth, call setup, and signalling functions, which are services that are extended to the ethernet-attached end stations desiring ATM services. The Figure also illustrates the cell switching data path, how proxy (virtually extended) services are extended to the ethernet-attached end stations, via the ATM Mgmt. & Signalling block.

GLOSSARY

10BASEF	Three standards for Ethernet/802.3 over fiber-optics. Specified in 802.3(TBD).
10BASET	Ethernet/802.3 over unshielded twisted pair. Specified in 802.3 (i).
802.1(d) Transparent Spanning Tree	This is the IEEE standard protocol for routing, loop detection and avoidance in a network of MAC bridges.
802.2 Logical Link Control	This is the IEEE standard for the Data Link Layer of local area networks. Also ISO 8802/2.
802.3	CSMA/CD LAN. This is the IEEE standardized LAN that is almost identical to Ethernet. It is a bus that uses CSMA/CD for access control. Also standardized as ISO 8802/3.
802.5 Token Ring	The IEEE standard for a ring using token passing for access control. Also known as Token Ring and standardized as IEEE 802.5 and ISO 8802/5.
AAL ATM Adaptation Layer	A protocol used to convert non-ATM information to and from ATM Cells.
ATM Asynchronous Transfer Mode	A switching/transmission technology which employs 53 byte cells as a basic unit of transfer. The ATM Cell is divided into 5 bytes of ATM Layer overhead and 48 bytes of ATM payload. ATM is fundamentally statistical in nature, with many "virtual circuits" sharing bandwidth.
ATM Adaptation	A protocol used to convert non-ATM information

SUBSTITUTE SHEET

- 41 -

Layer.	to and from ATM Cells.
CPSS Control Packet Switching System	A network protocol for communication between equipment. CPSS is a connection-less, packet-switched protocol used to transfer control and status information between network elements and network management entities.
CRC Cyclic Redundancy Check	A type of frame-check sequence.
DV Router Distance Vector Router	The Distance Vector (aka Bellman-Ford) routing algorithm used by CPSS. The only routing algorithm provided in CPSS-1.
DVMRP Distance Vector Multicast Routing Protocol	A routing table computation protocol for IP multicast services.
EGP Exterior Gateway Protocol	An IETF standard protocol used to propagate routing information between Autonomous Systems.
Ethernet	A CSMA/CD local area network developed at Xerox PARC during the late 70's. It is almost the same as the LAN standardized by the IEEE as 802.3.
FCS Frame Check Sequence	A checksum routine used to determine errors in the transmission of data packets.
FDDI Fibre Distributed Data Interface	An ANSI-standardized 100 Mbits/s local area network. The topology is a ring and uses token passing for access control.
HDLC High Level Data Link Control	A family of bit-oriented protocols providing frames of information with address, control and frame check sequence fields.
HEC	Header Error Check. The processing of validating the CRC-8 that protects the contents of the first four bytes of each ATM cell.
LAN Local Area Network	A system designed to inter-connect computing devices over a restricted geographical area (usually a mile or so).
LAT Local Area Terminal Protocol	A DEC-proprietary for terminal-host connections over a Local Area Network optimized for low-delay, high-bandwidths paths.
LATM Local ATM	A physical ATM link.
LLC Logical Link Control	See 802.2 above.
LS Router Link-State Router	The link-state (aka shortest path first) routing-table computation algorithm used by CPSS.
LUE Look Up Engine	A hardware table searching machine.
MAC	Media Access Control. The lower sub-layer of the Data Link Layer as described in the ISO Reference Model. The purpose of the MAC layer is to provide a reliable data transfer mechanism across a physical medium.

SUBSTITUTE SHEET

- 42 -

NIC Network Interface Card	An adapter board that can be added, usually by the user, to a workstation or PC in order to provided connectivity to a network, e.g., LAN or ATM.
NLSP NetWare Link Services Protocol	Novell's link-state (aka shortest path first) routing protocol for NetWare.
NMTI Node Management Terminal Interface	A local user interface for a network element.
NNI Network Node Interface	The interface between two network nodes which operate under different administrative domains.
OAM Cell Operations and Maintenance Cell	A specially tagged ATM cell (e.g. the header is different from the header for a user data ATM cell). OAM cells are specified to support ATM network maintenance features like Connectivity Verification, Alarm Surveillance, Continuity Check, and Performance Monitoring.
OSPF Open Shortest Path First	An IETF standard link-state routing protocol used for route determination in IP networks.
PCI Peripheral Component Interconnect	An emerging high-speed expansion bus for personal computers.
PTS Problem Tracking System	A centralized problem tracking database system used by Newbridge to manage the problems found in the field as well as within Newbridge (Design, PI, AppEng, etc.).
PVC Permanent Virtual Channel	An end-to-end logical ATM connection, of either the Virtual Channel or Virtual Path kind, established through administrative actions.
SNMP Simple Network Management Protocol	A standard for the management of entities in a TCP/IP local area network. There are two versions: the original (called SNMPv1) and a newer extended version (called SNMPv2).
SAR	<i>Segmentation and Reassembly</i> . The process of breaking non-ATM offered payloads into ATM cells (segmentation) and reconstructing ATM cells back into the circuit's native format (reassembly)
SVC	<i>Switched Virtual Channel</i> . A virtual channel that is established through signaling.
TCP	<i>Transmission Control Protocol</i> . The Transport Layer (roughly) protocol for the TCP/IP protocol suite. Defined in RFC 793.
Token Ring	See 802.5 above.
TXC	<i>Transmit Controller</i> . A ridge RISC complex for Ethernet transmission
UDP User Datagram Protocol	An unconfirmed datagram protocol used in IP networks. Defined in RFC 768.
UNI User-Network Interface	The interface between ATM user equipment and an ATM network.
VC Virtual Channel	A communication channel that provides for the

SUBSTITUTE SHEET

- 43 -

	sequential unidirectional transport of ATM cells.
VCC <i>Virtual Channel Connection</i>	A virtual channel that has end-to-end significance and is a concatenation of the virtual channel links that extends between the points where the ATM service users access the ATM Layer. The points at which the ATM cell payload is passed to, or received from, the users of the ATM Layer for processing signify the endpoints of a VCC.
VCI	<i>Virtual Circuit Identifier</i> . A 16 bit address used to identify an ATM connection.
VP	<i>Virtual Path</i> A logical association or bundle of VCs.
VPC <i>Virtual Path Connection</i>	A concatenation of virtual path links.
VPI <i>Virtual Path Identifier</i>	A 8 bit value used to identify an ATM path and carried in the cell header. Like a VCI it is locally significant and refers to the VPL active on the local UNI which comprises a large scope VPC.
WAN <i>Wide Area Network</i>	A system designed to interconnect computers within a campus.
XID <i>Exchange Identification</i>	A frame used for identification and Negotiation by many HDLC-like protocols.

SUBSTITUTE SHEET

- 44 -

Claims:

1. A computer communications network comprising a network of interconnected ATM switches over which ATM cells are transmitted, and a plurality of user devices
5 including LAN interface adapters for connection to one or more local area networks (LANs), characterized in that interface means connect at least some of said respective ATM switches to said LAN interface adapters, said interface means adapting said ATM cells for transport
10 over said LANS whereby said user devices can communicate through said LAN interface adapters transparently with said network.
2. A computer communications network as claimed in claim 1, characterized in that said interface means
15 comprises means for encapsulating said ATM cells in LAN frames for delivery directly to the LAN interface adapters of said user devices.
3. A computer communications network as claimed in claim 2, characterized in that said LAN interface
20 adapters are Ethernet adapters, and said encapsulating means encapsulates said ATM cells in Ethernet frames for delivery to said Ethernet adapters.
4. A computer communications network as claimed in claim 1, characterized in that said interface means
25 provide bridging, network-layer forwarding and LAN emulation functions to permit transparent communication between any of said user devices.

SUBSTITUTE SHEET

- 45 -

5. A computer communications network as claimed in claim 1, characterized in that it further comprises a network manager connected to at least one of said ATM switches, said network manager permitting an operator
5 thereof to configure said network.

6. A computer communications network as claimed in claim 1, characterized in that it further comprises route server means connected to one of said ATM network switches, said route servier means containing
10 admistrative information defining a virtual LAN including said ATM network.

7. A computer communications network as claimed in claim 1, further comprising monitoring means for permitting remote monitoring of said interface means,
15 said remote monitoring means transmitting data including errors from a monitored port to a remote monitoring port to replicate the data from said monitored port at said remote monitoring port, whereby tests can be conducted on said monitored port remotely as if on-site.

20 8. A computer communications network as claimed in claim 4, characterized in that it further comprises means for tagging received monitored packets to prevent said packets being treated as normally received packets by the receiving interface means.

25 9. A computer communications network as claimed in claim 1, characterized in that it comprises means for exchanging tokens between said interface means and said user devices, whereby a predetermined amount of data is

SUBSTITUTE SHEET

- 46 -

sent by a device receiving a token before the token is returned to the interface means, which then sends out a predetermined amount of data before passing on the token.

10. A device for interfacing an ATM computer
5 communications network with local area network (LAN) adapters, characterized in that it comprises a plurality of ports for connection to respective local area network adapters; means for receiving incoming LAN frames at any one of said ports; ATM segmentation means for creating
10 outgoing ATM cells; means for transmitting the outgoing ATM cells over an ATM network; means for receiving incoming ATM cells from the ATM network; re-assembly means for re-assembling LAN frames from incoming ATM cells; and frame queuing and transmission means for
15 transmitting outgoing LAN frames over a selected said port.

11. A device as claimed in claim 10, characterized in that it further comprises bypass means for directing local LAN traffic directly from said LAN receiving means
20 to said LAN frame queuing and transmission means, whereby said device can act as a bridge.

12. A device as claimed in claim 11, characterized in that it further comprises buffer means for buffering incoming and outgoing LAN frames on said ports, said
25 buffer means including means for reading MAC headers in incoming LAN frames.

13. A device as claimed in claim 11, characterized in that said buffer means output data to said ATM

SUBSTITUTE SHEET

- 47 -

segmentation means on n-bit wide buses, said device further comprising a translation engine for ensuring DMA packet transfer between said buffer means and said ATM segmentation means.

5 14. A device as claimed in claim 13, characterized in that said translation engine comprises means for discarding, bridging or network layer forwarding incoming LAN packets, means for encapsulating the packets for ATM layer adaptation, and means for inserting a virtual
10 channel.

15 15. A device as claimed in claim 14, characterized in that said translation comprises a RISC processor.

16. A device as claimed in claim 15, characterized in that said segmentation means comprises an ATMizer,
15 segmenter RAM on a primary bus thereof, high speed pointer memory on a secondary bus, and a dual-ported memory to said re-assembly means.

17. A device as claimed in claim 10, characterized in that it comprises means for passing tokens to said local
20 area networks, said tokens authorizing the transmission of data from a receiving host.

18. A device as claimed in claim 10, characterized in that port availability is provided as a bit mask derived from hardware signals and data availability is presented
25 as a bit mask in a packet address.

19. A device as claimed in claim 16, further comprising a priority encoder with round-robin priority for

- 48 -

returning a port with the next highest priority in the event of unservicibility of the current queue.

20. A device as claimed in claim 10, characterized in that it further comprises means for determining packet
5 length, a segmenter RAM for storing data packets to be segmented, a segmenter, a bus permitting DMA access of data packets in said segmenter RAM, and means for changing data on said data bus after a last byte has been read during DMA access.

10 21. A computer communications network characterized in that it comprises a plurality of interconnected ATM switches, a network manager and a route server connected to at least one of said ATM switches, a plurality of devices with LAN interface adapters, and means for
15 permitting said LAN interface adapters to communicate with said ATM switches.

22. A computer communications network as claimed in claim 21, characterized in that said permitting means comprises means for encapsulating said ATM cells in LAN
20 frames for delivery directly to said LAN interface adapters of some of said devices, and means for providing bridging, network-layer forwarding and LAN emulation functions to forward LAN frames to said LAN interface adapters of others of said devices.

SUBSTITUTE SHEET

1/10




FIG. 1






FIG. 2

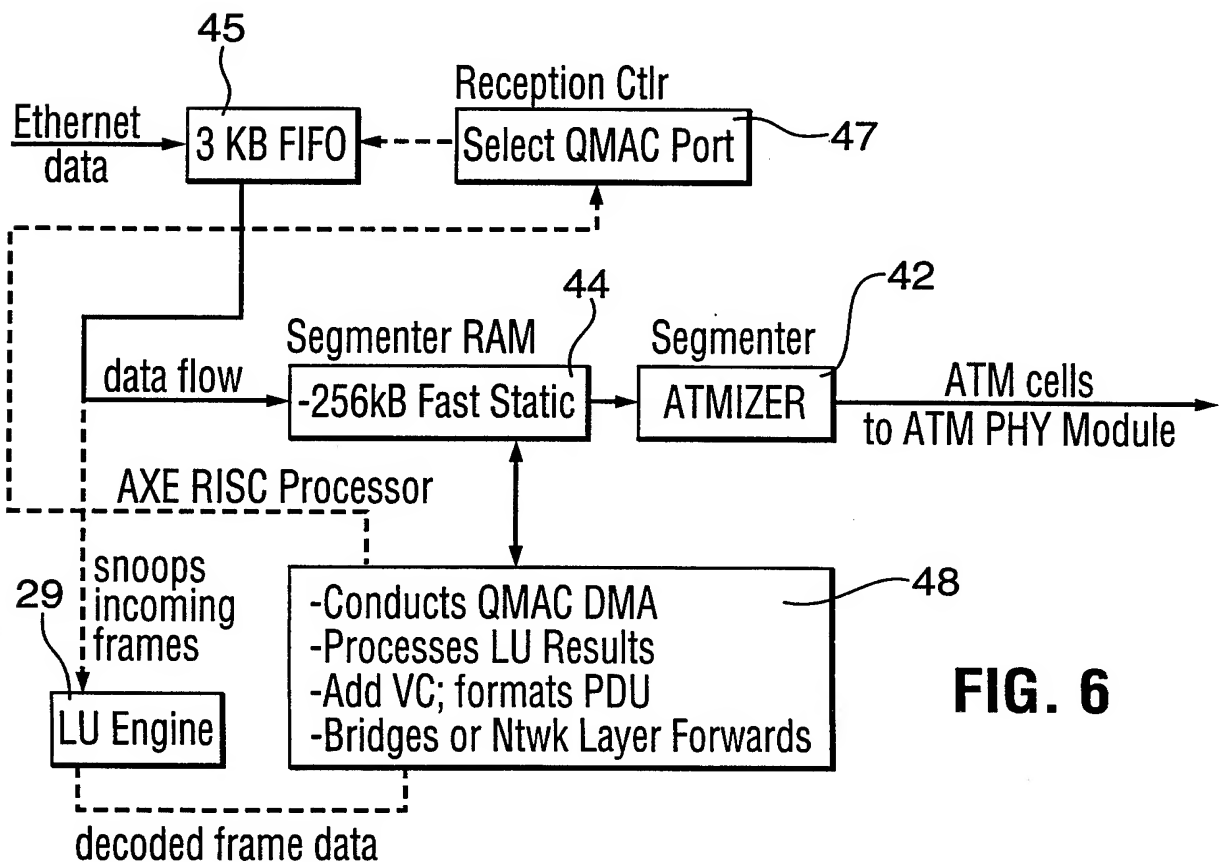
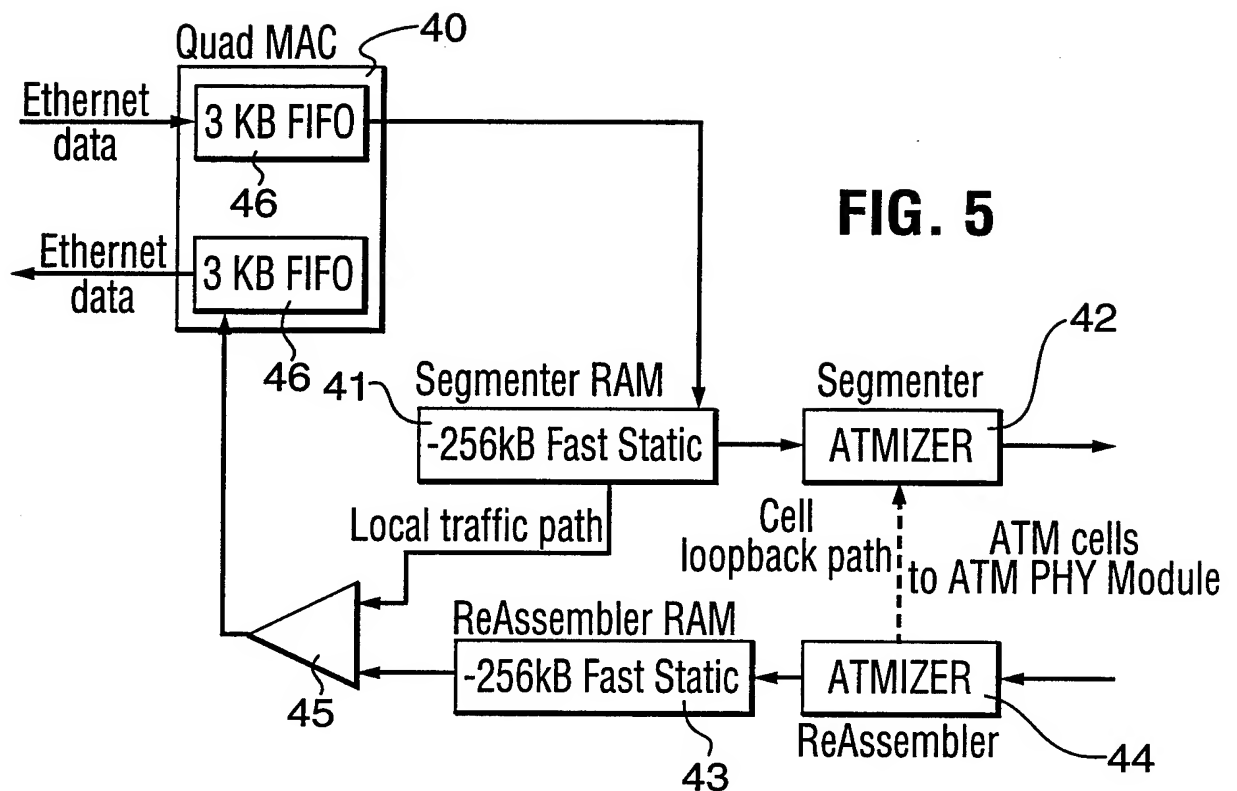
SUBSTITUTE SHEET

2/10



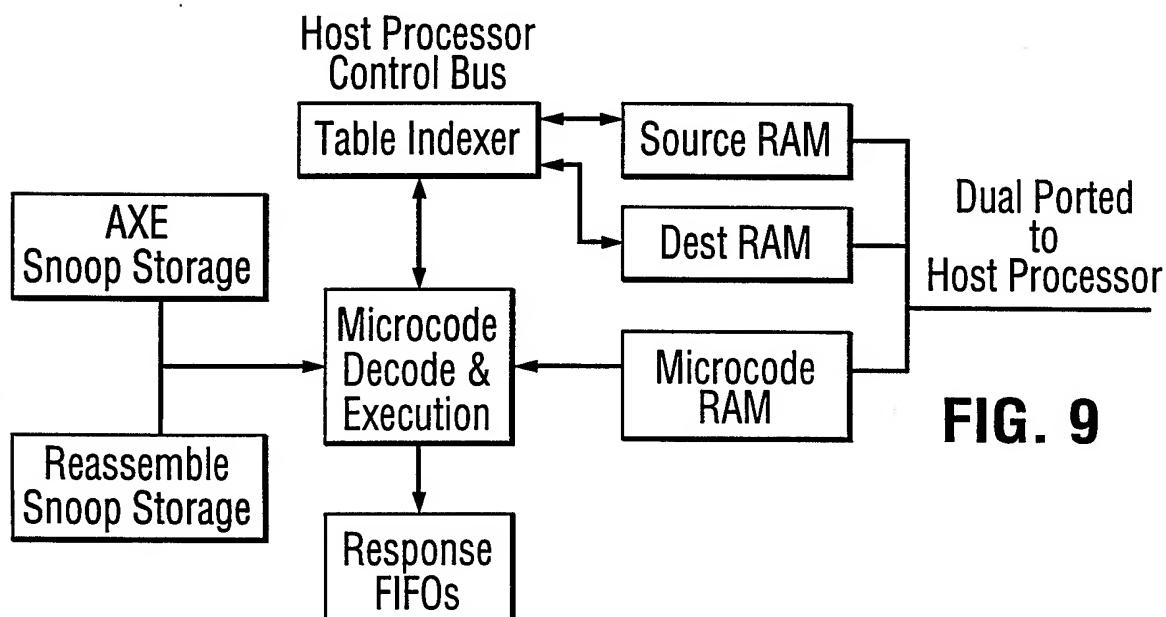
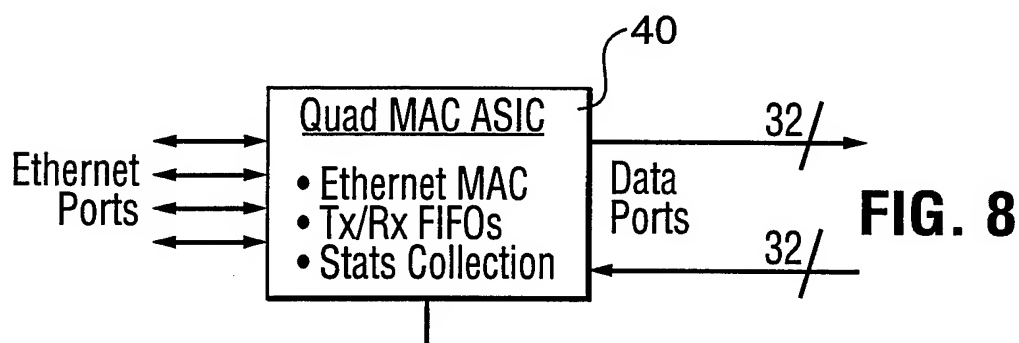
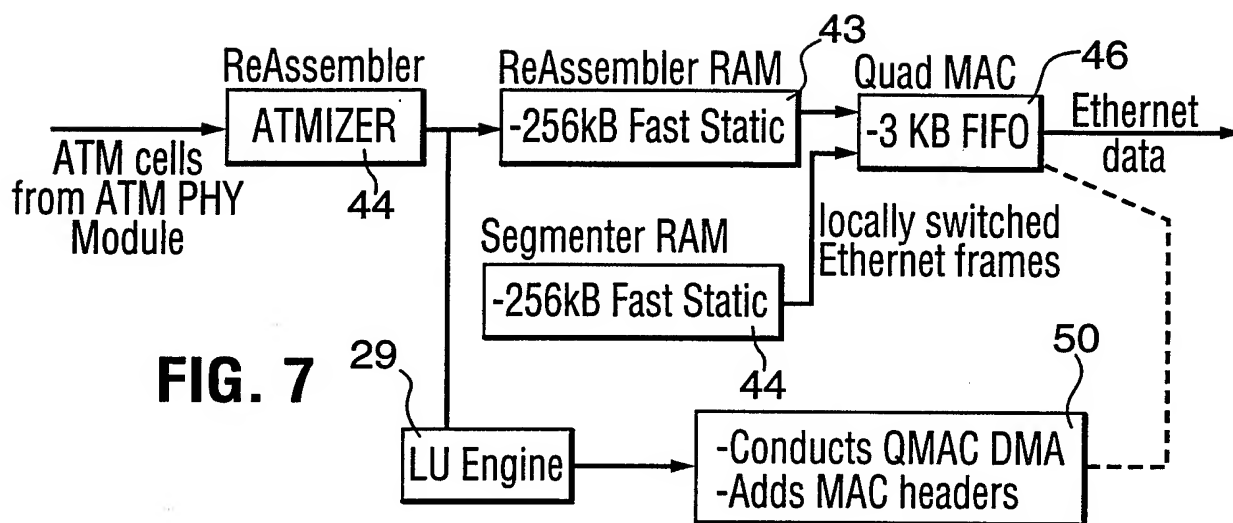
SUBSTITUTE SHEET

3 / 10






SUBSTITUTE SHEET

4 / 10



SUBSTITUTE SHEET

5 / 10



SUBSTITUTE SHEET

6 / 10

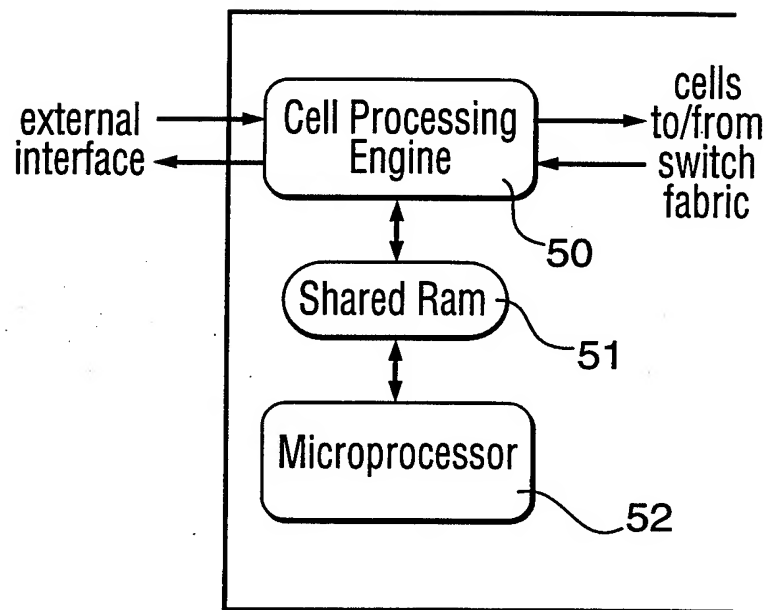


FIG. 13
PRIOR ART

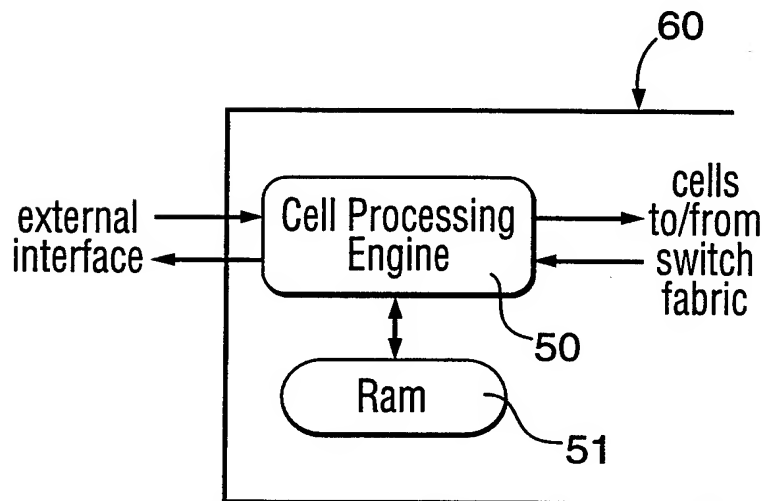



FIG. 14

7 /10

**FIG. 15**




FIG. 16

9 / 10

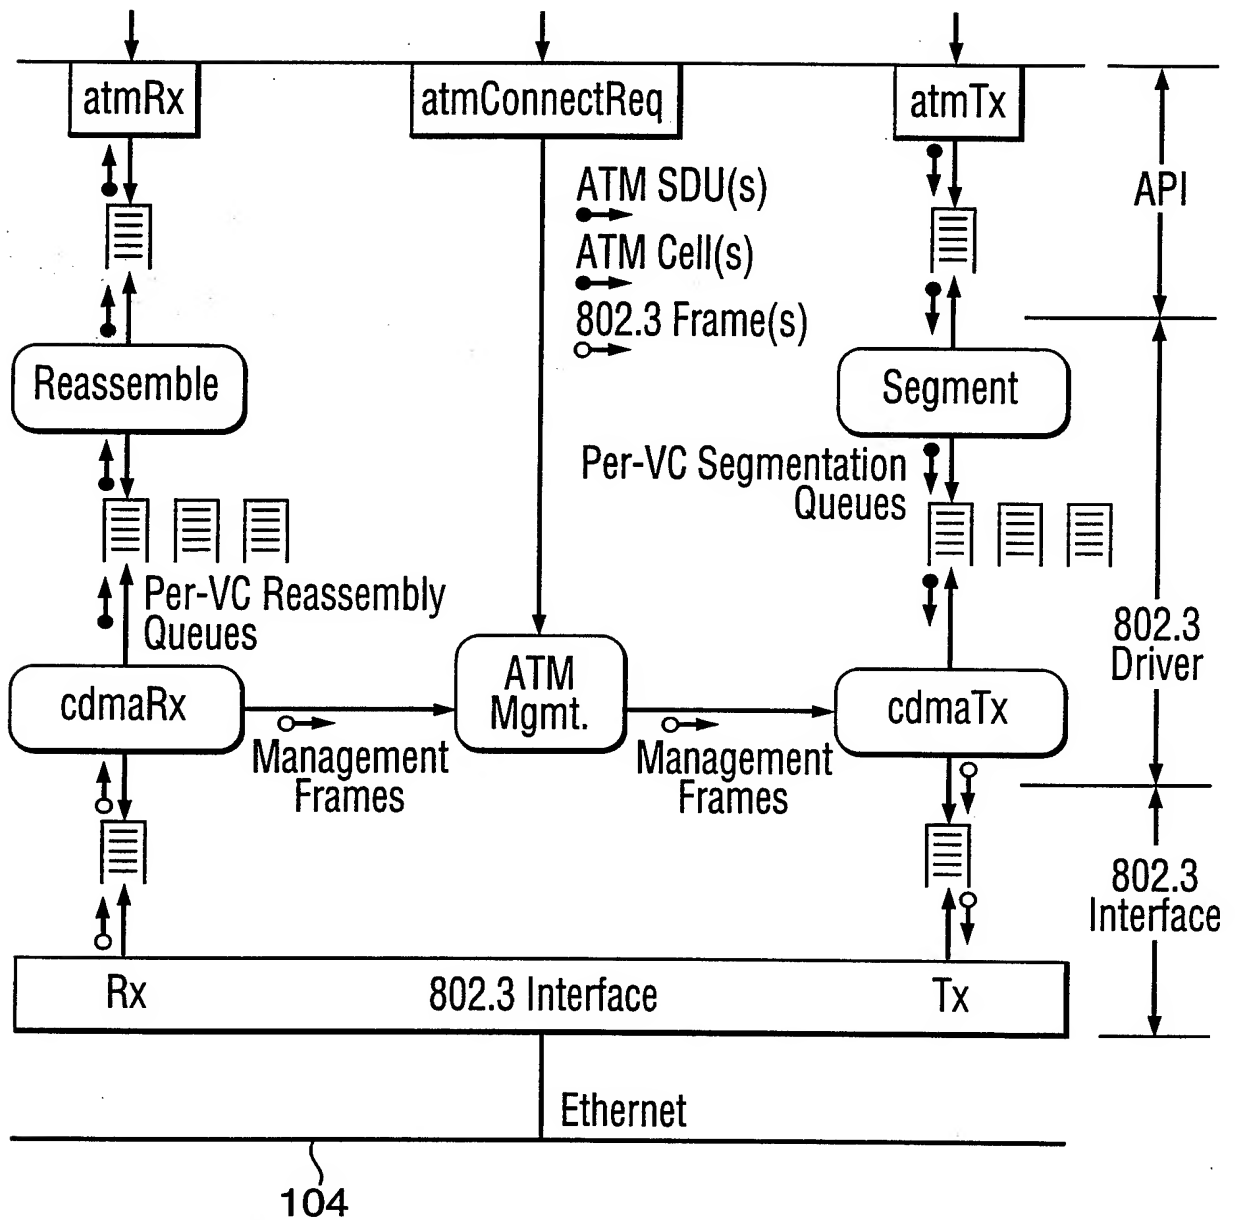
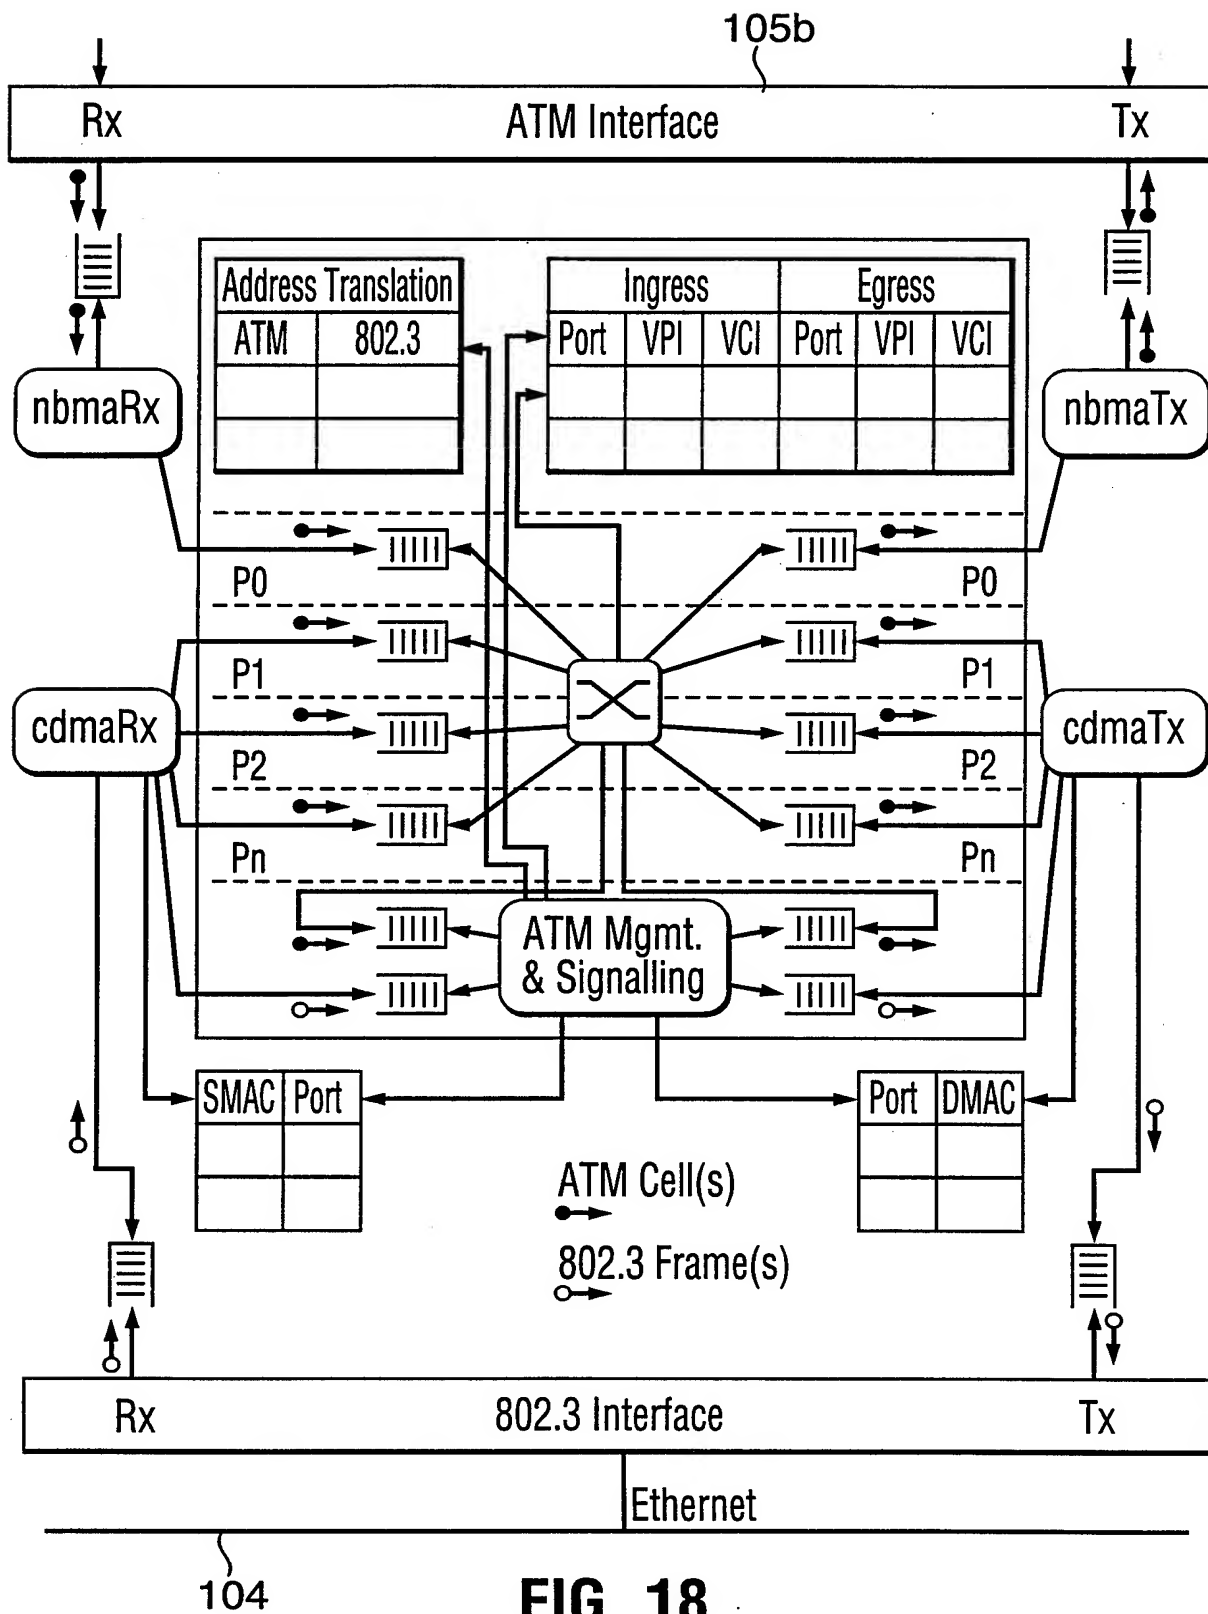


FIG. 17

10 / 10



INTERNATIONAL SEARCH REPORT

Intern. Appl. No.

PCT/CA 95/00029

A. CLASSIFICATION OF SUBJECT MATTER
 IPC 6 H04L12/46 H04Q11/04

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 H04L H04Q

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO,A,93 26107 (WASHINGTON UNIVERSITY) 23 December 1993 see page 2 - page 3; claims 1-4 see page 7, line 9 - line 28 ---	1-6, 10-13, 21,22
X	EP,A,0 524 316 (FUJITSU LIMITED) 27 January 1993 see column 4, line 35 - column 6, line 24 ---	1,6,10
X	EP,A,0 473 066 (MITSUBISHI DENKI KABUSHIKI KAISHA) 4 March 1992 see column 1, line 37 - column 2, line 3 see column 3, line 9 - line 50; claim 1; figure 9 --- -/--	1,10-12

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- "&" document member of the same patent family

Date of the actual completion of the international search

20 April 1995

Date of mailing of the international search report

0 4. 05. 95

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
 NL - 2280 HV Rijswijk
 Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
 Fax (+31-70) 340-3016

Authorized officer

Veen, G

INTERNATIONAL SEARCH REPORT

International Application No

PCT/CA 95/00029

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>PATENT ABSTRACTS OF JAPAN vol. 18 no. 38 (E-1495) ,20 January 1994 & JP,A,05 268256 (FUJITSU LTD) 15 October 1993, see abstract</p> <p style="text-align: center;">-----</p>	1,2

INTERNATIONAL SEARCH REPORT

Information on patent family members

Intern. Appl. No.

PCT/CA 95/00029

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO-A-9326107	23-12-93	AU-B- 4527793	04-01-94
EP-A-524316	27-01-93	AU-A- 1220892	07-09-92
		CA-A- 2079484	01-08-92
		WO-A- 9214321	20-08-92
EP-A-473066	04-03-92	JP-A- 4107029	08-04-92
		US-A- 5329527	12-07-94